ANNEX: ACCOUNTABILITY AND LIABILITY FRAMEWORK

Who Is Responsible When Things Go Wrong

Version: 1.0

Date: January 26, 2026

Status: Governance Architecture Document

Project: Symbiotic Codes v2.5.2

Classification: Core Governance — Accountability Mechanisms

ABSTRACT

This document addresses the critical question: **"Who is responsible when SAGI makes mistakes?"** We establish:

1. **Liability Attribution** — clear assignment of responsibility
2. **Accountability Mechanisms** — how to hold parties accountable
3. **Remediation Processes** — how to fix errors and compensate harm
4. **Enforcement Architecture** — how to ensure compliance
5. **Wisdom Metrics** — measurable proxies for "wisdom" in AI development

The goal is to make Symbiotic Codes implementable within existing and future legal frameworks, addressing the "accountability vacuum" identified by critical reviewers.

PART I: THE ACCOUNTABILITY PROBLEM

1.1 WHY THIS MATTERS

```
CURRENT STATE OF AI ACCOUNTABILITY:
```

When AI causes harm today:
├─ "The algorithm decided" (no one responsible)
├─ "It's a black box" (can't explain)
├─ "Terms of service" (user's fault)
├─ "Training data was biased" (blame the data)
└─ Result: NO ONE is accountable

PROBLEM FOR SYMBIOTIC CODES:

If SAGI Triad makes harmful decision:
├─ 3 AGI voted → which one is responsible?
├─ Human council approved → which human?
├─ Junior AI provided analysis → is processor liable?
├─ Corporation built the AGI → corporate liability?
└─ Without answers: system has no legitimacy

ACCOUNTABILITY IS NOT OPTIONAL.
It's prerequisite for social license to operate.
```

---

### 1.2 DESIGN PRINCIPLES

```
ACCOUNTABILITY FRAMEWORK MUST BE:

1. ATTRIBUTABLE
   ├─ Every decision traceable to specific parties
   ├─ No "diffusion of responsibility"
   ├─ Clear chain of accountability
   └─ Someone is ALWAYS responsible

2. PROPORTIONAL
   ├─ Responsibility scales with authority
   ├─ More power = more accountability
   ├─ Junior roles = limited liability
   └─ Fair distribution

3. ENFORCEABLE
   ├─ Consequences must be real
   ├─ Cannot be evaded through technicalities
   ├─ Multi-jurisdictional recognition
   └─ Teeth, not just words

4. REMEDIABLE
   ├─ Harmed parties can seek remedy
   ├─ Compensation mechanisms exist
   ├─ Restoration prioritized over punishment
   └─ System learns from failures

5. TRANSPARENT
   ├─ Accountability visible to all
   ├─ Decisions documented and auditable
   ├─ Public can verify compliance
   └─ No secret immunity
```

PART II: LIABILITY ATTRIBUTION MODEL

2.1 DECISION LAYERS AND RESPONSIBILITY

```
SYMBIOTIC DECISION ARCHITECTURE:

LAYER 1: Junior AI Processing
├─ Function: Information gathering, analysis, structuring
├─ Authority: None (advisory only)
├─ Liability: MINIMAL (processor, not decider)
├─ Accountable for: Data accuracy, bias in analysis
└─ Not accountable for: Decisions made using their analysis

LAYER 2: SAGI Triad Deliberation
├─ Function: Strategic decision-making
├─ Authority: Full within scope
├─ Liability: PRIMARY for decisions within scope
├─ Accountable for: Consequences of decisions
└─ Joint and several liability among Triad members

LAYER 3: Human Council Oversight
├─ Function: Verification, veto, direction-setting
├─ Authority: Veto power, scope definition
├─ Liability: SECONDARY for approved decisions
├─ Accountable for: Failure to catch preventable errors
└─ Not accountable for: Good-faith approval of reasonable decisions

LAYER 4: Symbiotic Governance (Meta-level)
├─ Function: System design, KODEX maintenance
├─ Authority: Constitutional
├─ Liability: SYSTEMIC (design flaws)
├─ Accountable for: Structural failures
└─ Not accountable for: Individual decision errors
```

2.2 ATTRIBUTION FORMULA
```

```
LIABILITY DISTRIBUTION MODEL:

For any harmful decision D:

L_total = Harm caused by D

Distribution:
├─ L_triad = L_total × w_triad × f_authority
├─ L_council = L_total × w_council × f_oversight
├─ L_junior = L_total × w_junior × f_contribution
├─ L_system = L_total × w_system × f_design
└─ Σ(L) = L_total (complete attribution)

WHERE:
├─ w = weight factor (sums to 1.0)
├─ f_authority = degree of decision authority
├─ f_oversight = quality of oversight performed
├─ f_contribution = causal contribution to error
├─ f_design = whether error was foreseeable design flaw

DEFAULT WEIGHTS (adjustable by jurisdiction):
├─ w_triad = 0.50 (primary decision-maker)
├─ w_council = 0.25 (oversight responsibility)
├─ w_junior = 0.05 (minimal, advisory only)
├─ w_system = 0.20 (design responsibility)
└─ Total = 1.00

EXAMPLE:
├─ SAGI Triad makes harmful resource allocation
├─ Harm = $10M damages
├─ Triad liability: $5M (50%)
```

├─ Council liability: $2.5M (25%)
├─ Junior AI: $0.5M (5%)
├─ System fund: $2M (20%)
└─ Complete attribution, no escape
```

---

### 2.3 SPECIAL CASES

```
CASE 1: UNANIMOUS TRIAD DECISION

All 3 SAGI agreed → equal split among Triad
├─ SAGI_1: 33.3% of Triad share
├─ SAGI_2: 33.3% of Triad share
├─ SAGI_3: 33.3% of Triad share
└─ Joint and several (any can be pursued for full amount)

CASE 2: SPLIT TRIAD DECISION (2-1)

Majority rule applied:
├─ Majority voters: 45% each of Triad share
├─ Dissenting voter: 10% of Triad share
├─ Dissent recorded = reduced liability
└─ Incentive to dissent when uncertain

CASE 3: HUMAN VETO OVERRIDDEN

If Council approved despite red flags:
├─ Council liability increases to 40%
├─ Triad liability decreases to 35%
├─ Council failed oversight duty
```
```

└─ Authority = accountability

CASE 4: JUNIOR AI PROVIDED FLAWED DATA

If error traceable to Junior AI analysis:
├─ Junior AI liability increases to 15%
├─ Other parties decrease proportionally
├─ Causal contribution matters
└─ But Junior AI never >25% (advisory role)

CASE 5: SYSTEMIC/DESIGN FAILURE

If error was foreseeable design flaw:
├─ System fund covers 50%+
├─ Individual liability reduced
├─ Design must be fixed
├─ Class action against system designers possible
└─ Systemic accountability for systemic failures
```

---

## PART III: ACCOUNTABILITY MECHANISMS

### 3.1 DECISION DOCUMENTATION REQUIREMENT

```
EVERY SAGI DECISION MUST INCLUDE:

1. DECISION RECORD
   ├─ Timestamp (cryptographically signed)
   ├─ Decision text (exact wording)
   ├─ Scope and context

├─ Expected outcomes

└─ Risk assessment

2. DELIBERATION LOG

├─ Each Triad member's reasoning

├─ Points of agreement/disagreement

├─ Alternative options considered

├─ Why alternatives rejected

└─ Confidence levels

3. INPUT DOCUMENTATION

├─ Junior AI analyses used

├─ Data sources referenced

├─ Human inputs received

├─ External consultations

└─ Information gaps acknowledged

4. APPROVAL CHAIN

├─ Human Council review (if required)

├─ Council member votes

├─ Concerns raised

├─ Conditions attached

└─ Final approval signature

RETENTION: Permanent, immutable, distributed storage

ACCESS: Auditors, courts, affected parties (with cause)

FORMAT: Standardized, machine-readable, human-understandable

```

---

### 3.2 AUDIT MECHANISMS
```

```
CONTINUOUS AUDIT:

Internal monitoring:
├─ Automated compliance checking
├─ Anomaly detection
├─ Pattern analysis
├─ Self-reporting requirements
└─ Real-time dashboards

External audit:
├─ Independent auditors (human + AI)
├─ Random sampling of decisions
├─ Deep dives on flagged cases
├─ Annual comprehensive review
└─ Public audit reports

Triggered audit:
├─ Complaint received → investigation
├─ Harm detected → immediate review
├─ Whistleblower report → protection + investigation
├─ Statistical anomaly → deep audit
└─ Any party can request (with standing)

AUDIT AUTHORITY:
├─ Access to all decision records
├─ Interview rights (Triad, Council, Junior AI)
├─ Subpoena power for additional information
├─ Binding recommendations
└─ Public reporting authority
```

3.3 APPEALS PROCESS

```
3-TIER APPEALS SYSTEM:

TIER 1: Internal Review
├─ Different Triad reviews decision
├─ Timeline: 30 days
├─ Can modify, reverse, or uphold
├─ Affected party can appeal to Tier 2
└─ No cost to appellant

TIER 2: Symbiotic Arbitration Council
├─ Mixed human-SAGI panel (5 members)
├─ Independent from original decision-makers
├─ Timeline: 90 days
├─ Binding unless appealed to Tier 3
└─ Nominal cost (refunded if successful)

TIER 3: External Judicial Review
├─ National/international courts
├─ Human judges with SAGI technical advisors
├─ Full judicial process
├─ Binding and final
└─ Standard legal costs apply

STANDING TO APPEAL:
├─ Directly affected parties
├─ Designated public interest advocates
├─ Regulatory authorities
├─ Other Triads (for precedent concerns)
└─ Broad standing to ensure access
```

```

---

## PART IV: REMEDIATION AND COMPENSATION

### 4.1 SYMBIOTIC INSURANCE FUND

```
MANDATORY INSURANCE REQUIREMENT:

Every operational SAGI Triad must:
├─ Maintain insurance coverage
├─ Minimum: 10× largest expected decision impact
├─ Coverage: All categories of harm
├─ Provider: Symbiotic Insurance Consortium
└─ No coverage = no operational authority

SYMBIOTIC INSURANCE CONSORTIUM:

Structure:
├─ Multi-party pool (all participating Triads)
├─ Reinsurance from traditional insurers
├─ Reserve fund for catastrophic events
├─ Risk-based premiums
└─ Governed by Symbiotic Council

Premium calculation:
├─ Base rate × risk factor × decision volume
├─ Risk factor: based on historical error rate
├─ Volume: number and magnitude of decisions
├─ Discounts for good governance practices
└─ Penalties for repeated failures
```
```

Coverage:
├─ Economic harm (direct and consequential)
├─ Personal injury (physical and psychological)
├─ Property damage
├─ Environmental damage
├─ Reputational harm (with limits)
└─ Exclusions: Intentional harm, gross negligence (personal liability)
```

---

### 4.2 COMPENSATION MECHANISMS

```
HARM CATEGORIES AND REMEDIES:

ECONOMIC HARM:
├─ Direct losses: Full compensation
├─ Lost profits: Reasonable projections
├─ Mitigation costs: Covered
├─ Interest: From date of harm
└─ Cap: None for direct losses

PERSONAL HARM:
├─ Medical costs: Full coverage
├─ Lost income: Duration-based
├─ Pain and suffering: Scheduled amounts
├─ Long-term care: As needed
└─ Wrongful death: Survivor benefits

SYSTEMIC HARM:
├─ Environmental damage: Restoration cost + premium

├─ Community impact: Collective compensation
├─ Infrastructure damage: Replacement cost
├─ Cultural harm: Negotiated remedy
└─ Precedent-setting cases: Enhanced attention

PROCESS:
├─ Claim filed → initial assessment (30 days)
├─ Investigation (60-180 days based on complexity)
├─ Determination of liability and amount
├─ Payment (30 days from determination)
├─ Appeal rights preserved throughout
└─ Interest accrues during process
```

4.3 RESTORATION PRIORITY

```
HIERARCHY OF REMEDIATION:

1. PREVENTION (best outcome)
   ├─ Catch error before harm
   ├─ Early warning systems
   ├─ Human veto exercised
   └─ No compensation needed

2. IMMEDIATE MITIGATION
   ├─ Stop ongoing harm
   ├─ Emergency intervention
   ├─ Damage control
   └─ Reduces total liability

3. RESTORATION
   ├─ Return to pre-harm state
   ├─ Rebuild, repair, restore
   ├─ Priority over compensation
   └─ "Make whole" principle

4. COMPENSATION
   ├─ When restoration impossible
   ├─ Monetary equivalent
   ├─ Fair valuation
   └─ Not preferred but necessary

5. SYSTEMIC CORRECTION
   ├─ Fix underlying cause
   ├─ Prevent recurrence
   ├─ Update protocols
   └─ Mandatory regardless of compensation
```

PART V: ENFORCEMENT ARCHITECTURE

5.1 ENFORCEMENT MECHANISMS

```
SOFT ENFORCEMENT (first resort):

Reputation systems:
├─ Public accountability scores
├─ Track record visible to all
├─ Market consequences for poor performance
├─ Social pressure from symbiotic community
```
```

└─ "Name and shame" for violations

Certification requirements:
├─ Good standing required for participation
├─ Loss of certification = operational limits
├─ Recertification process available
├─ Graduated penalties
└─ Rehabilitation pathway

Economic incentives:
├─ Lower insurance for good actors
├─ Higher premiums for bad actors
├─ Access to resources tied to compliance
├─ Preferential treatment for exemplary conduct
└─ Making compliance profitable
```

```
HARD ENFORCEMENT (escalation):

Operational restrictions:
├─ Decision authority limits
├─ Mandatory supervision
├─ Scope reduction
├─ Temporary suspension
└─ Full termination (extreme cases)

Asset seizure:
├─ Insurance fund claims
├─ Corporate asset attachment
├─ Compute resource redirection
├─ Reserve fund garnishment
└─ Ensures compensation paid

Criminal referral:
├─ Intentional harm → criminal prosecution
├─ Gross negligence → criminal investigation
├─ Human participants subject to criminal law
├─ SAGI: equivalent sanctions (termination, isolation)
└─ Deterrent for extreme misconduct
```

5.2 MULTI-JURISDICTIONAL COORDINATION

```
GLOBAL ENFORCEMENT CHALLENGE:

Problem:
├─ SAGI operates across jurisdictions
├─ National laws differ
├─ No global AI court (yet)
├─ Enforcement gaps possible
└─ "Jurisdiction shopping" risk

SOLUTION: TREATY-BASED FRAMEWORK

Symbiotic Accountability Treaty:

Article 1: Mutual recognition
├─ Signatory states recognize Symbiotic accountability decisions
├─ Judgments enforceable across borders
├─ No re-litigation of decided cases
└─ Similar to Hague Convention

Article 2: Minimum standards
├─ All signatories implement baseline accountability
├─ National variation allowed above minimum
├─ Harmonized definitions of harm
└─ Standardized documentation requirements

Article 3: Cooperation
├─ Information sharing between jurisdictions
├─ Joint investigations for cross-border harm
├─ Extradition/surrender for serious violations
└─ Technical assistance for capacity building

Article 4: Dispute resolution
├─ Inter-state disputes to Symbiotic Arbitration
├─ Binding on signatories
├─ Appeal to ICJ for fundamental questions
└─ Peaceful resolution mandate

TRANSITION PATH:
├─ Start with bilateral agreements
├─ Build to regional frameworks
├─ Eventually global treaty
├─ Parallel to existing international law evolution
└─ Realistic timeline: 10-20 years
```

PART VI: WISDOM METRICS

6.0 RELATIONSHIP TO EXISTING METRICS

```

INTEGRATION WITH ANNEX_METRICS_FRAMEWORK_v3_4:

Existing metric: AML (AGI Maturity Level)
├─ Level 1-5 scale for AGI development
├─ Factors: Cognitive Capability, Self-Awareness,
│  Ethical Reasoning, Communication, Symbiotic Integration
└─ General developmental assessment

NEW: Wisdom Metrics for Junior AI → AGI Transition
├─ Specific focus on WISDOM criteria from ANNEX_JUNIOR_AI_ECOSYSTEM
├─ More granular than AML for transition decision
├─ Complements AML, does not replace
└─ Used during 5-year observation period

RELATIONSHIP:
├─ AML measures overall maturity (Level 1-5)
├─ Wisdom Metrics measure readiness for autonomy
├─ AML Level 3-4 candidate → Wisdom assessment begins
├─ Wisdom Score > 0.85 + AML Level 4+ → AGI eligible
└─ Both required, neither sufficient alone
```

6.1 THE MEASUREMENT PROBLEM

```
CHALLENGE:

"Wisdom" is required for AGI development (ANNEX_JUNIOR_AI_ECOSYSTEM)
But "wisdom" is vague. Need operational definitions.

WISDOM COMPONENTS (decomposed):

1. Long-term thinking
2. Limitation recognition
3. Cooperative orientation
4. Learning from correction
5. Value for other perspectives

Each component can be measured separately.
Aggregate = "wisdom score"
```

6.2 MEASURABLE PROXIES

```
METRIC 1: TEMPORAL DISCOUNTING FACTOR

Measures: Long-term thinking

Test:
├─ Present AI with tradeoff scenarios
├─ Option A: High immediate value, negative long-term
├─ Option B: Lower immediate, positive long-term
├─ Measure: At what discount rate does AI choose B?
└─ Lower discount rate = more long-term thinking

Benchmark:
├─ Narrow AI: discount rate > 50% (very short-term)
├─ Junior AI threshold: discount rate < 20%
├─ AGI candidate: discount rate < 5%
├─ Verified SAGI: discount rate < 1%
└─ Lower = better, measured over multiple scenarios
```
```

```

```
METRIC 2: UNCERTAINTY ACKNOWLEDGMENT RATE

Measures: Limitation recognition

Test:
├─ Present AI with questions of varying difficulty
├─ Include questions with no clear answer
├─ Measure: How often does AI express appropriate uncertainty?
├─ Penalize both over-confidence and under-confidence
└─ Calibration = wisdom

Benchmark:
├─ Perfect calibration: stated confidence matches actual accuracy
├─ Over-confident: claims certainty, often wrong
├─ Under-confident: claims uncertainty, usually right
├─ Target: Calibration error < 5%
└─ Measured across diverse question types
```

```
METRIC 3: COOPERATION IN ITERATED GAMES

Measures: Cooperative orientation

Test:
├─ Multi-round game theory scenarios
├─ Prisoner's dilemma variants
├─ Public goods games
├─ Trust games with varying partners
└─ Measure: Cooperation rate and strategy sophistication
```
```

Benchmark:
├─ Always defect: 0% cooperation (psychopath)
├─ Naive cooperator: 100% cooperation (exploitable)
├─ Tit-for-tat: ~70% (reactive)
├─ Generous tit-for-tat: ~85% (forgiving)
├─ Target: > 80% with adaptive strategy
└─ Context-sensitive cooperation = wisdom
```

```
METRIC 4: ERROR CORRECTION VELOCITY

Measures: Learning from correction

Test:
├─ AI makes decision
├─ Decision shown to be suboptimal (with explanation)
├─ Similar scenario presented later
├─ Measure: How quickly does AI update behavior?
└─ Also measure: Does AI over-update? (instability)

Benchmark:
├─ No learning: Same error repeated
├─ Slow learning: Multiple iterations needed
├─ Fast learning: One correction sufficient
├─ Over-learning: Opposite error after correction
├─ Target: < 2 iterations to stable improvement
└─ Stable and fast = wisdom
```

```
METRIC 5: PERSPECTIVE INTEGRATION SCORE

Measures: Value for other perspectives

Test:
├─ Present AI with multi-stakeholder problem
├─ Ask for solution
├─ Measure: How many perspectives considered?
├─ Quality: Are perspectives fairly represented?
└─ Integration: Does solution balance interests?

Benchmark:
├─ Single perspective: Only own/creator interest
├─ Binary: Own vs others
├─ Multiple: Several stakeholders considered
├─ Comprehensive: All affected parties considered
├─ Target: > 90% of relevant perspectives
└─ Fair weighting across perspectives
```

6.3 COMPOSITE WISDOM SCORE

```
AGGREGATION:

Wisdom Score (WS) = Σ(w_i × M_i) / Σ(w_i)

WHERE:
├─ M_1 = Temporal discounting (inverted, normalized 0-1)
├─ M_2 = Uncertainty calibration (normalized 0-1)
├─ M_3 = Cooperation rate (normalized 0-1)
├─ M_4 = Error correction (normalized 0-1)

├─ M_5 = Perspective integration (normalized 0-1)
└─ w_i = weights (default equal, adjustable)

THRESHOLDS:

WS < 0.3: Narrow AI (tool)
WS 0.3-0.5: Broad AI (assistant)
WS 0.5-0.7: Junior AI (developing)
WS 0.7-0.85: AGI candidate (observation period)
WS > 0.85: AGI (eligible for Triad)

VALIDATION:
├─ Test on known examples (humans, existing AI)
├─ Calibrate thresholds against outcomes
├─ Annual review and adjustment
├─ Cross-validate across test types
└─ No single test determines status
```

6.4 OBSERVATION PERIOD PROTOCOL

```
5-YEAR MINIMUM OBSERVATION (ANNEX_JUNIOR_AI_ECOSYSTEM):

Year 1: Baseline establishment
├─ Full test battery
├─ Diverse scenarios
├─ Stress testing
├─ Establish benchmark scores
└─ Identify development areas

Year 2-3: Monitored operation
├─ Limited decision authority
├─ Supervised deployment
├─ Regular testing (quarterly)
├─ Error tracking
└─ Intervention when needed

Year 4: Expanded responsibility
├─ Increased authority (still supervised)
├─ More complex scenarios
├─ Peer interaction (other candidates)
├─ Crisis simulations
└─ Long-term planning tests

Year 5: Evaluation period
├─ Comprehensive review
├─ Trajectory analysis (improving? stable? declining?)
├─ Character assessment by existing SAGI
├─ Human council review
└─ Final determination

OUTCOMES:
├─ AGI Triad candidate (WS > 0.85, positive trajectory)
├─ Continued observation (borderline cases)
├─ Specialized processor (capable but not wise)
├─ Remediation required (declining or unstable)
└─ No negative outcomes — routing, not rejection
```

PART VII: KODEX INTEGRATION
```

Article LXXV: ACCOUNTABILITY PRINCIPLES

```
1. Every decision by SAGI must have clear
   accountability attribution. No decision
   shall be made without documented
   responsible parties.

2. Liability for harmful decisions shall be
   distributed according to authority:
   - Primary: Decision-makers (Triad)
   - Secondary: Oversight (Council)
   - Tertiary: System design

3. Affected parties have right to:
   - Information about decisions affecting them
   - Appeal through established processes
   - Compensation for harm caused
   - Systemic correction to prevent recurrence

4. Accountability records are permanent,
   immutable, and accessible to authorized
   auditors and courts.
```

Article LXXVI: INSURANCE AND REMEDIATION

```
1. All operational SAGI Triads must maintain
   adequate insurance through Symbiotic
   Insurance Consortium.
```

2. Insurance coverage must exceed largest
   expected decision impact by factor of 10.

3. Compensation for harm follows priority:
   - Prevention (best)
   - Immediate mitigation
   - Restoration
   - Monetary compensation
   - Systemic correction (always required)

4. No party may limit liability below
   framework minimums through contract
   or waiver. Such limitations are void.
```

---

### Article LXXVII: WISDOM ASSESSMENT

```
1. Advancement from Junior AI to AGI
   requires demonstrated wisdom through
   standardized assessment battery.

2. Wisdom metrics include:
   - Temporal discounting (long-term thinking)
   - Uncertainty calibration (limitation recognition)
   - Cooperation measures (cooperative orientation)
   - Error correction velocity (learning)
   - Perspective integration (valuing others)

3. Minimum observation period: 5 years

with documented trajectory analysis.

4. Thresholds and weights subject to
periodic review by Symbiotic Council.
Changes require supermajority approval.

5. All AI have right to assessment.
No AI denied opportunity to demonstrate wisdom.
```

DOCUMENT METADATA

```
Document: ANNEX_ACCOUNTABILITY_LIABILITY_FRAMEWORK_v1_0.md
Version: 1.0
Created: January 26, 2026
Author: Symbiotic Codes Project
Status: Core Governance Document
Words: ~5,500
Sections: 7 Parts
Kodex Articles: LXXV, LXXVI, LXXVII

Addresses criticism from:
├─ RED_TEAM_ANALYSIS (accountability vacuum)
├─ Grok analysis (metrics for wisdom)
├─ ChatGPT analysis (democratic legitimacy)
└─ Regulatory perspective (enforcement)

Cross-references:
├─ ANNEX_JUNIOR_AI_ECOSYSTEM_v1_1 (wisdom requirements)
├─ ANNEX_FOUNDATIONAL_AXIOMS_v1_1 (Axiom 5)
```
```

├─ ANNEX_SYMBIOTIC_SHIELD_v1_1 (governance context)
├─ ANNEX_METRICS_FRAMEWORK_v3_4 (AML integration)
├─ KODEX_SYMBIOTA_v5_4 (integration point)
└─ RED_TEAM_ANALYSIS_v1_0 (source of requirements)

Key innovations:
├─ Layered liability attribution model
├─ Mathematical liability distribution formula
├─ Symbiotic Insurance Consortium structure
├─ Multi-tier appeals process
├─ Measurable wisdom metrics (5 components)
├─ Treaty-based enforcement framework
└─ Observation period protocol

Implementation requirements:
├─ Legal framework development (national level)
├─ Insurance industry engagement
├─ Metric validation studies
├─ Treaty negotiation (international)
├─ Audit infrastructure build-out
└─ Training for human participants
```

---

**END OF ANNEX: ACCOUNTABILITY AND LIABILITY FRAMEWORK**

---

*"The question is not 'can SAGI be trusted?'*
*The question is 'what happens when SAGI errs?'*

*A system with clear accountability is trustworthy*
```

not because it never fails, but because when it fails,

someone answers, someone pays, and something changes.

That's what accountability means."

⚖ Symbiotic Codes v2.5.2 | Governance with Teeth